Combined Unscented Kalman and Particle Filtering for Tracking Closely Spaced Objects*

Robert J. Pawlak
Code T41, NSWC
Dahlgren, VA 22448
robert.pawlak@navy.mil

***Abstract* - Tracking closely spaced objects with resolution limited sensors is a difficult problem. One way to address this issue is to track these targets individually, and employ relatively complex data association approaches as a means of pairing detections and tracks. The algorithm outlined in this paper takes a different approach, and instead estimates the group velocity using an unscented Kalman Filter (UKF). The UKF state estimate is then employed within a particle filter, which estimates the distribution of objects within the group. It is shown that this approach can be very effective, especially for groups of irregularly spaced objects.**

Keywords: Tracking, merged measurements, multiple measurements, tracking, unscented kalman filter, particle filter, surface radar.

1 Introduction

In many instances, sensor performance can be limited in resolution and/or update rate. Under these conditions, the association of data to tracks can be extremely problematic, especially when there are multiple, closely spaced objects. Indeed, issues of limited power, resolution and update rate are very important when considered in the context of networked sensor grids.

Thus, in addition to the normal problems experienced with intermittent target detection, tracking algorithms also need to deal with the so-called "merged measurement" [1] problem. This phenomenon occurs when multiple targets fall in the same resolution cell, producing only a single detection or return. And while merged measurements also present a problem, multiple measurements can also result from a single object, especially for slow sensor update rates. Indeed, the purpose of the algorithm presented in this paper is to deal with these two difficulties.

Various methods have been proposed to deal with these phenomena, including multiple hypothesis tracking and various probabilistic data association methods. The method discussed in this paper makes no attempt to maintain individual tracks, and instead tracks all objects within a group as a single entity. Once a state space estimate of the group behavior is obtained, a particle filter is employed to develop a super-resolution estimate of the distribution of objects within the group. This estimate can then be used to determine raid count, or to get a rough estimate of the distribution of objects within a group. These two pieces of information can be vital in the context of situational awareness, or for targeting purposes.

The motivation for the approach discussed in this paper was the problem of tracking large numbers of closely spaced boat targets with a surface search radar [2]. The algorithmic development assumes measurements in a polar coordinate system, with relatively poor azimuth resolution and slow update rate. The parameters chosen for the notional sensor do not correspond to any real sensor. However, they are commensurate with the resolution limitations one encounters when dealing with ubiquitous, inexpensive surface search radar systems.

Another unique aspect of this algorithm is that the state of the particle filter is itself a probability density function. This gives rise to a whole host of issues associated with evaluating the likelihood function, which is required for the particle filter weight updates.

The steps in the algorithm are as follows. First, the centroid position of the grouped measurements is used to update an unscented Kalman filter estimate of the Cartesian position, velocity and acceleration. This state estimate is then passed to a particle filter which is responsible for developing a super-resolution estimate of the target positions within the group. The state space of the particle filter is the infinite set of feasible two-dimensional target probability density functions. The particle filter updates are computed by using a two-dimensional correlation function to estimate the likelihood.

Note that the data for a particular group track are assumed to be associated by a prior, independent process. In this respect, the work in this paper differs significantly from other work, in which the data association process is explicitly integrated into the tracking algorithm [3][4]. For many surface search radars, this is done via a clustering technique with clustering neighborhood size sufficient to address worst case target dynamics and sensor resolution issues. Clustering techniques are also popular for reducing computational load as well [1]. Structuring the algorithm in this fashion allows it to be easily placed into an adjunct processor which can work in conjunction with existing clustering systems.

*This work was supported by the Office of Naval Research, Code ONR351, Mr. Gil Graff.

Report Documentation Page

Form Approved
OMB No. 0704-0188

1. REPORT DATE **JUL 2006**	2. REPORT TYPE	3. DATES COVERED **00-00-2006 to 00-00-2006**
4. TITLE AND SUBTITLE **Combined Unscented Kalman and Particle Filtering for Tracking Closely Spaced Objects**		5a. CONTRACT NUMBER 5b. GRANT NUMBER 5c. PROGRAM ELEMENT NUMBER
6. AUTHOR(S)		5d. PROJECT NUMBER 5e. TASK NUMBER 5f. WORK UNIT NUMBER
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) **Naval Surface Warfare Center (NSWC),Code T41,Dahlgren,VA,22448**		8. PERFORMING ORGANIZATION REPORT NUMBER
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)		10. SPONSOR/MONITOR’S ACRONYM(S) 11. SPONSOR/MONITOR’S REPORT NUMBER(S)

12. DISTRIBUTION/AVAILABILITY STATEMENT
Approved for public release; distribution unlimited

13. SUPPLEMENTARY NOTES
9th International Conference on Information Fusion, 10-13 July 2006, Florence, Italy. Sponsored by the International Society of Information Fusion (ISIF), Aerospace & Electronic Systems Society (AES), IEEE, ONR, ONR Global, Selex - Sistemi Integrati, Finmeccanica, BAE Systems, TNO, AFOSR’s European Office of Aerospace Research and Development, and the NATO Undersea Research Centre.

14. ABSTRACT
see report

15. SUBJECT TERMS

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT **unclassified**	b. ABSTRACT **unclassified**	c. THIS PAGE **unclassified**	**Same as Report (SAR)**	**6**	

2 Measurement Model

The sensor measurements were generated using a simulation of a single 2-D radar with a rotating antenna and fixed scan period, but nothing within the algorithm formulation would preclude the use of multiple sensors. The parameters of this sensor are shown in table 1:

Table 1, Sensor parameters used to generate target returns

Parameter	Value
Scan/sample period (ΔT)	2 sec
Beamwidth ($\Delta\theta$)	1.1 deg
Range bin size (ΔR)	120 m

Referring to Figure 1, if a target is located in the shaded area, it will be assigned a specific azimuth and range value. This sensor model more accurately simulates the noise characteristics (i.e. quantization) that is to be expected in an actual sensor, than the typically employed additive gaussian model[1]. In a radar sensor, the angular dimension $\Delta\theta$ is analogous to the 3dB beamwidth, and ΔR is the range bin size.

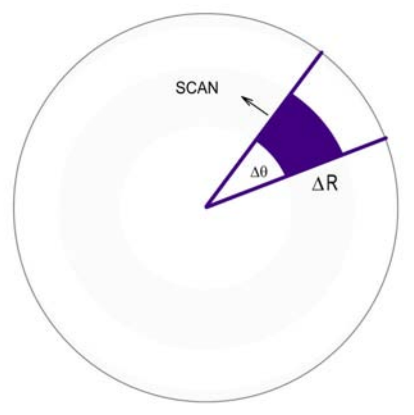

Figure 1: Sensor model

Let this quantization be represented by the following mapping:

$$z_k(r,\theta) = f(r',\theta') \tag{1}$$

where r' represents the radial distance to an object, and θ'represents the azimuth to the object, before quantization. The mapping is defined such that multiple objects in the same range and azimuth cell will result in only a single detection for that cell. Furthermore, as long as a target is within the beam (i.e. the beam is swept across an object and dwells upon it for a finite amount of time), target returns will result. This means that a single target may produce multiple detections. These sensor sampling and quantization effects are many times ignored when assessing track filter performance. As will be subsequently shown, they have a great effect on tracking performance, especially in the context of the group tracking problem.

3 Unscented Kalman Filter Implementation

The UKF formulation chosen herein relies on filtering an averaged group position. Proximate measurements associated with a single group and arriving on the same scan are averaged, and provided as a single measurement to the filter. If N_z measurements are used in the centroid operation, an approximation to the measurement noise variance in range and angular coordinates is given by: $\frac{\Delta R^2}{12N_z}$, and $\frac{\Delta\Theta^2}{12N_z}$, respectively[2]. The nonlinear measurement equations for the UKF are given by the following equations:

$$\overline{r}_k = \frac{1}{N_z}\sum_{l=1}^{N_z} r_k^l \tag{2}$$

$$\overline{\Theta}_k = \frac{1}{N_z}\sum_{l=1}^{N_z} \Theta_k^l \tag{3}$$

$$H(\overline{\Theta}) = atan(\frac{y}{x}) \tag{4}$$

$$H(\overline{r}) = \sqrt{x^2+y^2} \tag{5}$$

With the measurement vector defined as:

$$\overline{z}_k = \begin{bmatrix} \overline{r}_k \\ \overline{\Theta}_k \end{bmatrix} \tag{6}$$

and the measurement noise covariance approximated by:

$$R = \begin{bmatrix} \frac{\Delta R^2}{12N_z} & 0 \\ 0 & \frac{\Delta\Theta^2}{12N_z} \end{bmatrix} \tag{7}$$

Even though the applied part of this research is directed toward tracking groups of boats, a simple six state (employing two independent three-state models) dynamic model was chosen. While there are specific boat dynamic models (see [7] for an example), the group dynamics of the centroid of a number of boats will, in general, be quite different from that of a single object. Therefore, the state vector and state transition matrix are given by:

$$s_k = \begin{bmatrix} x_k \\ \dot{x}_k \\ \ddot{x}_k \\ y_k \\ \dot{y}_k \\ \ddot{y}_k \end{bmatrix} \tag{8}$$

[1] It is important to note that the UKF filter models employ a gaussian model for the measurement error, however, the *measurements* employed in the filter use the quantization noise discussed herein.

[2] This is the resulting variance of the average of N_z independent measurements.

and

$$\Phi = \begin{bmatrix} F & 0 \\ 0 & F \end{bmatrix} \tag{9}$$

where

$$F = \begin{bmatrix} 1 & \Delta T & \frac{\Delta T^2}{2} \\ 0 & 1 & \Delta T \\ 0 & 0 & 1 \end{bmatrix} \tag{10}$$

Where ΔT is the sample period between centroid estimates. The process noise covariance was defined as:

$$Q = \gamma \begin{bmatrix} q & 0 \\ 0 & q \end{bmatrix} \tag{11}$$

with a white noise model assumed:

$$q = \begin{bmatrix} \frac{\Delta T^5}{20} & \frac{\Delta T^4}{8} & \frac{\Delta T^3}{6} \\ \frac{\Delta T^4}{8} & \frac{\Delta T^3}{3} & \frac{\Delta T^2}{2} \\ \frac{\Delta T^3}{6} & \frac{\Delta T^2}{2} & \Delta T \end{bmatrix} \tag{12}$$

In equation (11), γ is a design parameter - for the results discussed in this paper, $\gamma = 0.05$. It should be noted that the centroid of the group of objects traveling together will have less dynamic capability than a single object.

Equations (2) - (12) were employed in the UKF formulation detailed in [8], along with values of $\alpha = 0.0001$, $\kappa = 0$, $\beta = 2$. A covariance expansion term of 1.05 was used at each iteration of the filter to avoid divergence of the estimate. The latter is very likely to occur in the event of an intra-group crossing maneuver after the filter estimates have settled.

4 Particle Filter Implementation

The intent of the particle filter is to estimate the probability density function of objects within the group track (in Cartesian coordinates). But with the current algorithm structure, the best that can be done is to estimate the probability density function (PDF) of the centroid, and hope that peaks within the PDF somehow correspond to object locations. This is not an unreasonable assumption.

Each particle within the filter is determined by its weight and support points. The number of support points for each particle is given by M, with N being the total number of particles. Each particle is comprised of a set of support points:

$$\phi_k^i(x, y) = \{(\overset{*}{x}{}^i_{j,k}, \overset{*}{y}{}^i_{j,k}),\ j = 1 \dots M\} \tag{13}$$

Where $\overset{*}{x}{}^i_{j,k}$ is the x coordinate of j^{th} support point of the i^{th} particle on the k^{th} iteration, $\overset{*}{y}{}^i_{j,k}$ is defined in an analogous fashion. Note that from this point forward, the dependence of ϕ_k^i on $(\overset{*}{x}, \overset{*}{y})$ is implicitly assumed.

This work uses the bootstrap particle filter formulation, in which the update step is accomplished via a straightforward propagation of the particles through the state space. This is accomplished by updating the support of each particle, under the assumption that each particle corresponds to a tracked object, that travels with approximately the same group velocity as determined by the centroid. Therefore, the support for each particle is updated via the mapping:

$$\phi_k^i \to \phi_{k+1}^i, i = 1 \dots N \tag{14}$$

This mapping is completely determined by the update of the support points in Cartesian coordinates:

$$\overset{*}{x}{}^i_{j,k+1} = \overset{*}{x}{}^i_{j,k} + \Delta T\, \dot{x}_k + \frac{\Delta T^2}{2}\, \ddot{x}_k + \epsilon \tag{15}$$

$$\overset{*}{y}{}^i_{j,k+1} = \overset{*}{y}{}^i_{j,k} + \Delta T\, \dot{y}_k + \frac{\Delta T^2}{2}\, \ddot{y}_k + \epsilon \tag{16}$$

Since the individual objects within the group do not always move in perfect synchrony with the centroid of the group, a small noise term (ϵ) is added to account for this uncompensated movement. The latter is assumed to be a $N(0, \sigma_p)$ distributed random variable. A relationship may also exist between this term and the covariance estimate developed with the UKF, but this relationship was not developed in the course of this research.

As can be seen from equations (13) - (16), each particle can be envisioned as a two-dimensional PDF with support points that shift with each filter update. The posterior density is given by the weighted sum of the particles:

$$p(\phi_k |\, z_{1..k}) = \sum_{i=1}^{N} w_k^i\, \delta\,(\phi_k - \phi_k^i) \tag{17}$$

With the particle weights defined by w_k^i, z_k is the set of sensor measurements $\{(r_k^l, \Theta_k^l),\ l = 1 \dots N_z\}$ for the current group track. Since the prior is used as the importance sampling distribution, the weight update equation becomes:

$$w_{k+1}^i = w_k^i\, p(z_k | \phi_k^i) \tag{18}$$

Because of the formulation chosen for this filter, an exact evaluation of $p(z_k|\phi_k^i)$ is difficult to implement. Therefore, it is helpful to think of the proposed distribution of object measurements resulting from a particular particle as a discrete PDF with support defined by equation (13).

The support points can then be transformed into polar coordinates via equations (4) and (5), and then quantized using the mapping defined in equation (1). The resultant distribution of target returns for a particular particle is now defined in polar coordinates, and is denoted by $\phi_k^i(\overset{*}{r}, \overset{*}{\theta})$, with support:

$$\phi_k^i = \{(\overset{*}{r}{}^i_{j,k}, \overset{*}{\theta}{}^i_{j,k}),\ j = 1 \dots M\} \tag{19}$$

The result can be envisioned as a PDF, which can then be compared to the PDF obtained from the actual sensor measurements. The most intuitive candidate for evaluating this likelihood is the chi-squared test. Unfortunately, one limitation of this test is its inability to deal with zero likelihood events, as well as events with a relatively small frequency of occurrence [5]. Another likely candidate is the Kolmogorov-Smirnov test. However in practice, this measure tends to be overly optimistic in its evaluation of the correspondence between two distributions, and thus produces extremely slow convergence to the proper weight values.

Given these difficulties, another goodness-of-fit measure was employed. Since each particle is a two-dimensional PDF, a straightforward visual analogy would suggest the use of the two-dimensional correlation function, which can be applied to the two-dimensional histogram, which is evaluated over the total feasible support of the particles (in measurement space) for the group track. Let $A_{m,n}$ denote the probability of occurrence (on iteration k, which is implicitly assumed from this point forward) for the proposed particle density at the m,n th histogram cell, computed for the support points given by (19). Likewise, $B_{m,n}$ denotes the probability density obtained from histogramming the measurements on iteration k (from equation (1)). The two-dimensional correlation of these two images will be given by:

$$\rho = \frac{\sum_{N_m}\sum_{N_n}(A_{m,n}-\overline{A})(B_{m,n}-\overline{B})}{\sqrt{\sum_{N_m}\sum_{N_n}(A_{m,n}-\overline{A})^2\sum_{N_\mathrm{m}}\sum_{N_n}(B_{m,n}-\overline{B})^2}} \tag{20}$$

Where N_m is the number of bins in the m dimension, N_n is the number of bins in the n dimension, $\overline{A}$ is the mean of matrix A and $\overline{B}$ is the mean of matrix B. This function is defined on $[-1,1]$, and hence a scaled version of the correlation is used, along with an acceleration term that speeds convergence of the weights.

Thus the likelihood is computed from the equation below:

$$p(z_k|\phi_k^i) \approx \left(\frac{\rho+1}{2}\right)^a \tag{21}$$

Where a is the aforementioned acceleration term, usually taken to be on the interval $[1,4]$. Substitution of (21) into (18) gives the weight update relationship.

As will be seen in subsequent sections, the combined UKF/Particle filter was applied in two different scenarios. In both cases, it will be shown that a performance increase results. However, the application of these filters requires that a number of design parameters be selected. The values for these parameters, as used in the particle filter are shown in table 2.

Table 2, Parameters used for the particle filter

Parameter	Value
σ_p	0.010
M	5
N	100
N_m	12
N_n	12
a	4

The standard assumption of uniform weights ($w_0^i = N^{-1}$, $\forall i$) is also used.

5 Scenarios

Two scenarios will be presented which show the potential advantage of applying the proposed hybrid tracking approach. One employs a group of targets moving in approximate synchrony, until separation and subsequent resolution of individual objects begins to occur at 250 sec. into the scenario. The other scenario has a group in which two group members execute a crossing maneuver. It is assumed that in both scenarios, all the objects are tracked as a single group, and that their close spacing, along with the resolution of the sensor, make this a necessity.

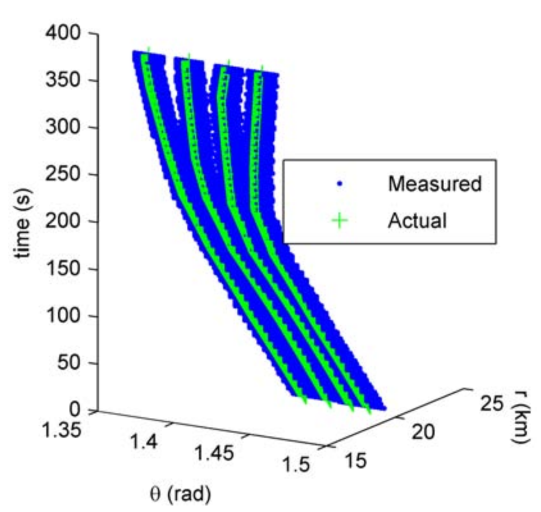

Figure 2: Four inbound objects scenario

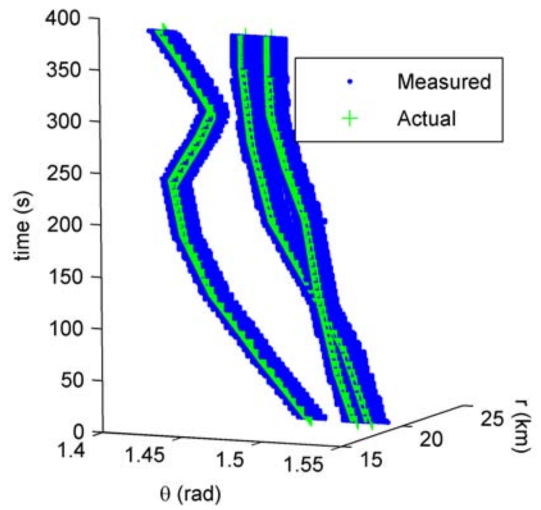

Figure 3: Three inbound w/two crossing objects

6 Findings

Clearly, it is difficult to synthesize the number of tracked objects from the particle filter estimate of the state[3]. However, it is possible to compare the estimate of the resultant target PDF to that obtained when using the centroid estimate (without any particle filtering). This was accomplished by computing the distance between the centroid/UKF estimates of the swarm position to the closest true object position. Likewise for the particle filter, the distance from the support point of the most probable object location (as defined by the maximum of the posterior - equation (17)) to the closest object position was computed. Henceforth, this number will be called the minimum distance (MD) metric. Plots of this error metric (in meters) for the two cases discussed in the previous section are shown in Figures 4 and 5.

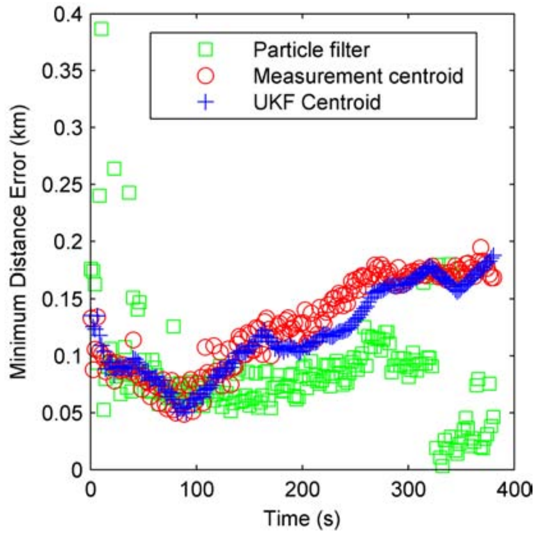

Figure 4: Performance, four inbound objects

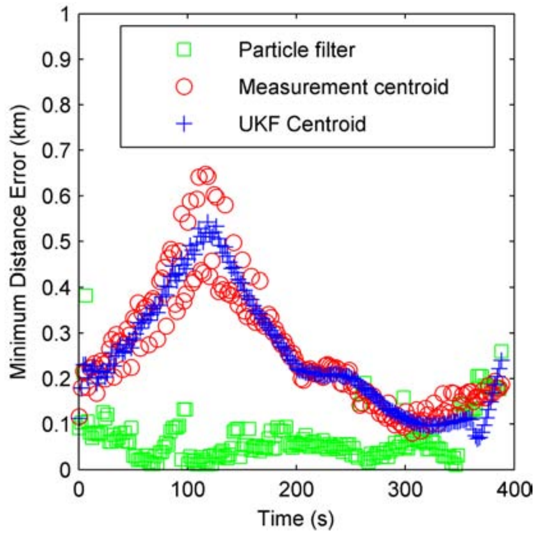

Figure 5: Performance, three inbound objects

[3] Such an estimate would require the examination of the two-dimensional PDF for local maxima. Accurate raid count would require an accurate mapping of these maxima, which is a difficult proposition.

The most performance increase is evident when the objects are irregularly distributed throughout the group and when there is some partial resolution of group members (figures 3 and 5). There is less payoff for applying the algorithm to large, regularly spaced target formations (figures 2 and 4). But note for this latter scenario, that there is even increased performance before the objects are partially resolved at 250 sec.

Note as well the relatively slow settling times of the filters. This is due in part to parameter choices, as well as the limited resolution and update rate of the notional sensor. Indeed, if this algorithmic approach has a drawback, it is certainly its relatively slow convergence time. However, it should also be noted that there were relatively few particles used to estimate the object positions (in this case, 100). This was due to the typical limitations of computer hardware, and it is expected that an increase in the number of particles, as well as other parameters associated with computing the likelihood (equation (21)) would enhance the performance of the hybrid filter.

Another finding is that the UKF only performs on par with the simple, unfiltered centroid estimate. This is chiefly due to resolution effects, and the relatively slow settling time of the filters (for an example of this, see figure 4). Table 3 summarizes the performance difference between the UKF and the centroid. It shows little or no performance gain from applying the UKF to obtain a better estimate of group position. However, what the UKF does provide (in the context of this algorithm), is an estimate of the group velocity, which is integral to the particle filtering estimate.

Table 3, Performance of UKF vs. Centroid Track

RMS Error in km	Four objects	Three objects
UKF (X)	0.02141	0.04672
Centroid (X)	0.01722	0.05174
UKF (Y)	0.00956	0.03616
Centroid (Y)	0.01010	0.02101

The mean and root mean square (RMS) of the MD metric was also computed for each scenario, and these errors are summarized in table 4. It's quite clear that there is a significant advantage accrued by adding the particle filter estimate (up to a factor of 4 improvement). And as to be expected in light of the previously discussed comparison of the UKF and centroid, there is relatively little performance difference when using the UKF to estimate a true object location versus using the centroid.

Table 4, Performance of Particle Filter, UKF, and Centroid as assessed by MD metric.

Mean/RMS error in km	Four objects	Three objects
PF	0.08403/0.04402	0.06880 /0.05296
UKF	0.11945/0.03724	0.25241/0.12932
Centroid	0.12676/0.04052	0.25927/0.13030

Note that a probability of detection of 1 was used to obtain these results. A lower probability of detection (and/or higher false alarm rate) would also adversely affect convergence time.

To summarize, the results indicate that the hybrid particle/UKF filtering approach provides improved performance over simply using the centroid. This is a valuable finding, because it suggests that the performance of many 2-D surface search radars[4] can be improved via the addition of an add-on approach similar to the one outlined in this paper. Such an approach would require minimal intervention into the front-end processing of the radar, and could work in conjunction with existing cluster-based trackers.

7 Conclusion

The hybrid UKF/Particle filter approach discussed in this paper is effective for improving the ability of a sensor to develop more accurate object position estimates in a group tracking framework. Within this context, the two largest issues are multiple detections per target, and merged detections for objects in the same resolution cell.

The hybrid UKF/particle approach showed up to a factor of 4 improvement in estimation error over the centroid, so there is a tangible benefit to using this approach to more accurately localize objects. The main drawbacks are its processing intensive nature, and the relatively long sample histories that are needed to develop the particle filtering estimate. Interestingly, there was relatively little performance gain associated with using the UKF versus using the centroid estimate.

It should also be noted that this approach would work well in the context of an adjunct processor, to improve object resolution in group and cluster based trackers. These trackers are used in many surface search radars.

And while it seems as though object count might be difficult to synthesize from the particle filter estimate, more work could be done in postprocessing the particle filter output to extract additional useful information.

References

[1] S. Blackman, R. Popoli, *Design and Analysis of Modern Tracking Systems*, Artech House, 1999.

[2] S. Horman, R. Pawlak, J. Bibel , C. Phillips, *Proposed Ship Self Defense Weapons System Design to Counter the Swarming Asymmetric Threat*, Proc 2004 National Fire Control Symposium, Aug 2004.

[3] H. Shyu, Y.Lin, J. Yang, H. Jinchi, *The Group Tracking of Targets on Sea Surface by 2-D Search Radar*, IEEE Int'l Radar Conf. pp. 329-333, May 1995.

[4] X. Wang and D. Musicki, *Low Elevation Sea-Surface Target Tracking Using IPDA Type Filters*, Proceedings Radar 2003, International Conference, pp. 472-478, Sept 2003.

[5] W. Conover, *Practical Non Parametric Statistics, Second Edition*, J. Wiley and Sons, 1980.

[6] G. Veres and J. Norton, *Improved Particle Filter for Multitarget-Multisensor Tracking with Unresolved Observations*, 2001 IEE Conf On Target Tracking: Algorithms and Applications, pp. 12/1-12/5 Vol. 1, 2005.

[7] E. Semerdjiev, L. Mihaylova, *Variable- and Fixed-Structure Augmented Interacting Multiple-Model Algorithms for Maneuvering Ship Tracking Based on New Ship Models*. International Journal of Applied Math Computer. Science, Vol.10, No.3, 591-604, 2000.

[8] R. van der Merwe and E. Wan, *Sigma-Point Kalman Filters for Probabilistic Inference in Dynamic State-Space Models,* Proc Workshop on Advances in Machine Learning, (Montreal, Canada), June 2003.

[4] Many of which employ clustering techniques for group tracking.